ENABLING HIGH PERFORMANCE RECONFIGURABLE GROUND DATA PROCESSING SYSTEMS

Marco Aurelio de Figueiredo
Stingher Gaffarian Technologies, Inc.
NASA Goddard Space Flight Center
Greenbelt, Maryland 20771
marco@fpga.gsfc.nasa.gov

Tonjua M. Hines
NASA Goddard Space Flight Center
Greenbelt, Maryland 20771
tonjua.hines@gsfc.nasa.gov

Abstract

The deployment of high data rate instruments onboard satellites with direct broadcast capability demands the implementation of powerful ground data processing systems. Conversely, to enable the widespread usage of direct broadcast data and the generation of subsequent data products, low cost data processing systems must be available. Reconfigurable computers based on Field Programmable Gate Arrays (FPGA) enable the implementation of ground data processing systems that lead to higher performance and lower cost while keeping a high degree of programmability and maintainability. This paper presents the conceptualization of the MODIS Adaptive Level One Accelerator (ALOA), a prototype under construction at NASA's Goddard Space Flight Center. The ALOA is a computer workstation augmented by a reconfigurable computer. The system will calibrate and geographically locate data from the EOS-AM1 MODerate resolution Imaging Spectroradiometer (MODIS) instrument at near real-time data rates.

Key words: EOS, MODIS, FPGA, reconfigurable computing, direct broadcast.

Introduction

The remote sensing user community demands ever-quicker access to the data collected by instruments onboard Earth observing satellites. The reduction of time is accomplished through the direct downlink of the science data to the user as the satellite passes over the user location on Earth. The science user can then process the collected data and produce data products at real time data rates. Examples of such products of interest are the determination of the perimeter of a forest fire or the area of an oil slick on a body of water, and the location and tracking of the eye of severe storms.

The above scenario is feasible today at reasonable cost only for low data rate instruments. Multispectral imaging instruments required to produce the products given as examples utilize data rates in the tens of mega bits per second. To preprocess the formatted spacecraft data for such instruments and apply the instrument calibration functions, powerful ground data processing systems are required. Even more processing power is required from these ground data systems to generate the end-user product. The cost associated with creating such a system to process the instrument data and create products at real-time data rates makes the solution prohibitive to most interested users. A technology that could increase the processing performance while reducing the cost of direct readout systems is thus welcome.

Adaptive computing, also known as reconfigurable computing, is an emerging technology that utilizes Field Programmable Gate Arrays (FPGA) to implement computationally intensive algorithms at the hardware gate level. Since a particular hardware architecture is implemented for each specific application, acceleration rates of several orders of magnitude faster than current computers are attainable. Furthermore, even though the application design process still requires a traditional hardware design approach, the end product of a hardware reconfigurable system is the software that is used to configure the FPGA devices. This software product, better defined as configware, can be modified to support system changes such as algorithm adjustments required throughout the life of a remote sensing instrument. Two implementations of remote sensing applications using adaptive computing have been developed by the Adaptive Scientific Data Processing group at Goddard Space Flight Center[6,7].

Both yield an order of magnitude acceleration over high-end workstations.

Accelerator cards utilizing arrays of FPGA devices organized in a configurable architecture are commercially available for both desktop and mini-computer systems. The Adaptive Level One Accelerator (ALOA) is a prototype under development that will process the MODIS Level 1 functions at near real-time data rates. The ALOA will utilize several adaptive accelerator cards to augment the computation capability of a desktop computer system. The expected result is a system affordable to those remote sensing users that require real-time data products.

This paper gives an overview of the MODIS direct broadcast system and the ground system composition. It then exposes the plan for the ALOA prototype system.

The EOS-AM1 Telemetry Systems

The AM-1 is the first of a series of satellites that will compose the Earth Observing System (EOS). The AM-1 telemetry systems can link to earth via the Tracking and Data Relay Satellite System (TDRSS) or by direct transmission to ground stations.

The EOS Data and Information System (EOSDIS) receives the AM-1 data stream from TDRSS. The data is first processed by the EOS Data and Operations System (EDOS) and sent by network to Distributed Active Archive Centers, or DAACs, where it is stored, processed, and made available to the user community by the EOS Core System (ECS). ECS processing consists of many steps, each resulting in a data product designed for a specific purpose. A data product may be distributed directly to the user community, used as input for later processing steps, or both.

The AM-1 Direct Access System (DAS) provides real-time science data from the spacecraft instruments directly to the science community independent of the EOSDIS[1]. The DAS provides three types of science data return capabilities: the real-time direct downlink (DDL) of ASTER instrument data, the real-time direct broadcast (DB) of MODIS instrument data and spacecraft ancillary data, and the direct playback (DP) link for a backup to the TDRSS Ku-band science data downlink. The DAS transmits at 8.2125 GHz (X-band) to user gound stations, which are supposed to have either a 3 meter or 11.3 meter receiver dish antenna. The 3 meter dishes are suitable only for the direct broadcast service data rates (13.125 Mbps). The DDL and DP modes require a larger dish to accommodate the 105 Mbps data rate.

The user must schedule DAS services for a time when the spacecraft passes over the user's ground station. The scheduling of DAS services will be achieved by Relative Time Sequences (RTS) commands uplinked from the EOS Operations Center (EOC) the day before the scheduled contact.

The MODIS Data Processing Flow

The MODerate resolution Imaging Spectroradiometer (MODIS), the flagship instrument aboard the AM-1 spacecraft, images the earth's atmosphere and surface in visual and infrared wavelengths with 12-bit precision[8]. Data is collected for 36 spectral bands at three resolutions: 1000, 500, and 250 meters/pixel. MODIS peak data production rate is approximately 11Mbits/sec, which corresponds to the day-time data collection when both the reflective and emissive bands produce valid data.

According to the EOSDIS data processing environment, a total of 48 data products organized in 5 (0 to 4) distinct levels are defined in the MODIS processing data flow. Figure 1 shows a high level view of the science data processing flow. Level 0 is processed by the EDOS' Level Zero Processing Facility (LZPF). Its primary function is to restore the data to the state in which it was recorded onboard the spacecraft. The level zero processing system sorts and time stamps the data packets, eliminating duplicate packets and including quality annotation. Further data processing is performed at the DAACs.

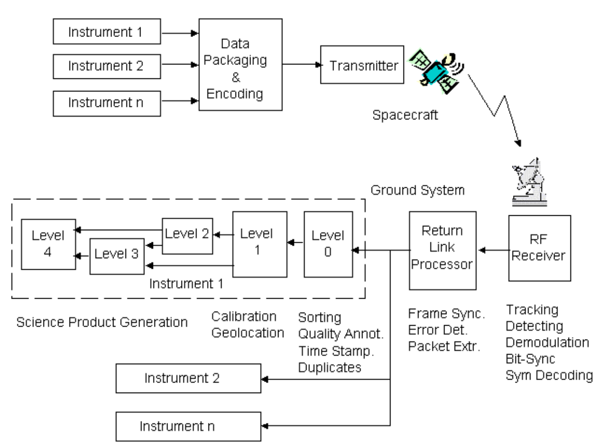

Figure 1 Science data processing functional diagram

The MODIS Level 1 processing data flow is composed of 3 functional blocks that produce three sets of data products (MOD01, MOD02, and MOD03) as illustrated in figure 2. The Level 0 data files contain the raw data stream from the MODIS instrument. "Raw data" includes digital representations of MODIS's sensor voltages, which in turn correspond to a scene of the earth's atmosphere and surface. Engineering data, such as instrument temperature, optical system parameters, calibration information, and geographical position are also present in Level 0 data.

The Level 1A data set includes all MODIS data in digitized (counts) form for all bands, all spatial resolutions, all time tags (converted), all detector views (Earth, solar diffuser, SRCA, black body, and space view), and all engineering and ancillary data, stored in HDF format, as are all subsequent data forms used in MODIS data processing. The geolocation data set contains geodetic coordinates, ground elevation, and solar and satellite zenith, and azimuth angle for each MODIS 1-km sample.

The Level 1B data set contains calibrated and geolocated at-aperture radiances for all spectral bands plus additional data including quality flags, error estimates and on-board calibration data. The Level 1B data is the input data used to generate every remaining MODIS product - it is used by theoretical models that use radiance as a starting point to make useful statements concerning atmospheric and geophysical properties.

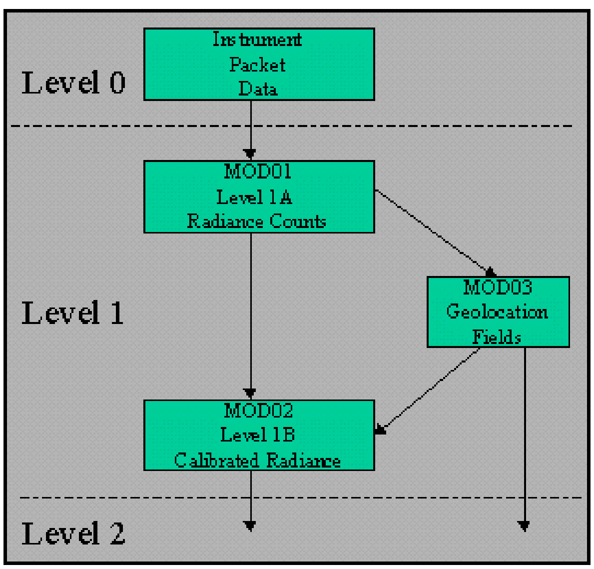

Figure 2 MODIS Level 1data flow diagram

A MODIS direct readout system is a rather simpler ground data processing system. Since only the MODIS data set is received on the AM-1 direct broadcast channel, there is no need to separate packets from different instruments. As a result , the Return Link Processing (RLP) and Level 0 functions are simplified. Certain characteristics of the Level 1 processing stages are also different from the EOSDIS processing data flow. For example, the largest data set that a direct readout system will receive is when the satellite crosses exactly above the ground station. This data set is equivalent to approximately 14 minutes of MODIS data. As a result, the calibration algorithms must be adjusted to comply with the direct readout characteristics. It is also expected that the EOSDIS requirement to store intermediary data in Hierarchical Data Format (HDF) be flexed to permit the data processing at real time data rates. Only the output of the direct readout system (level 1) will be stored as HDF.

The Adaptive Level One Accelerator (ALOA)

The MODIS direct readout system is expected to be composed of two workstation based processing systems. The 3-meter antenna, RF receiver, return link processor (RLP), and level zero processing elements will compose the data ingest system. This system will use commercial hardware components including a desktop PC running Microsoft Windows NT. The second processing stage will produce MODIS level 1 products. The later system is defined as the Adaptive Level One Accelerator (ALOA).

There are two configurations planned for the ALOA. The high end ALOA system will be composed of a dual 180MHz MIPS R10000 processors SGI Origin 200 workstation, augmented by 4 AMS Wildforce adaptive accelerators. The accelerator cards will be housed on a separate PCI extender chassis to provide an improved electrical and thermal environment. The lower cost system will be composed of a 400MHz Pentium based PC augmented by the same Wildforce cards used in the Origin 200, also housed on a PCI extension chassis.

The Level 1 software developed for the DAACs uses the ANSI C programming language. This software is being modified to adjust to the requirements of the direct readout system. The new software is being written using the Java programming language. Java is a modern object oriented language designed for distributed data processing systems. The language is also interpreted making it platform independent. Tests run by the ASDP group indicated that Java's just in time compiler reduces the performance advantage of a non-interpreted language such as C to a minimum. Besides, since the computation intensive portions of the code are going to be performed on the adaptive accelerators, the possible advantage of a pre-compiled code are literally non existent.

Furthermore, an object oriented programming environment provides an easier path to integration with the hardware oriented adaptive accelerators. Several trade-offs can be made where software objects are migrated to and from the adaptive accelerator to come up with the best performance arrangement.

The adoption of Java will also facilitate the transition of the accelerator cards from the SGI to the PC environment. No modification to the host code will be required. Benchmarking the two systems will be as simple as measuring the output data production rate of each. It is expected that the PC environment may provide the same computation performance while reducing the system cost by an order of magnitude.

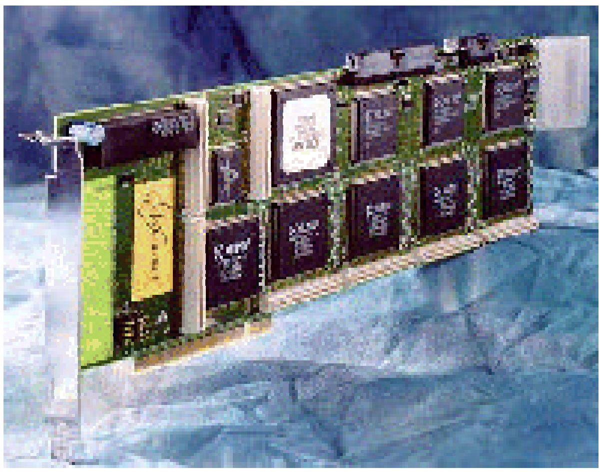

Figure 3 AMS Wildforce

The Annapolis Micro System (AMS) Wildforce WAC4085 adaptive accelerator is composed of 5 Xilinx 4085XLA FPGA devices organized as a systolic array (figure 3). Each FPGA device has a mezzanine connector that allows the addition of a dual-port SRAM memory module of up to 4 Mbytes. Together they compose what is called a processing element (PE). The dual-port memory allows direct communication between the host computer and the individual PEs through the PCI bus.

An FPGA based crossbar switch interconnects the 5 processing elements via a 36-bit channel allowing dynamic, host controlled, inter-PE connection (figure 4). A 36-bit systolic bus connects PE1 to PE4. The systolic bus can be extended to other Wildforce cards through IO connectors. The host can send and receive data from the systolic bus through bi-directional FIFOs connected to PE1 and PE4. A third bi-directional FIFO connects the host to PE0, also called the controlling processing element (CPE).

The ALOA system will contain four WAC4085 Wildforce cards housed on the host computer's PCI bus. A total of 1,700,000 equivalent hardware logic gates will be available for acceleration of the MODIS level 1

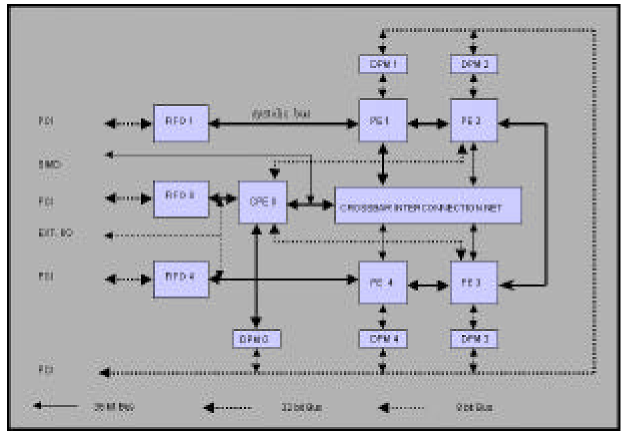

Figure 4 Wildforce Architecture

algorithms. The cards can be combined in several architectures. For example, the same card can be used to accelerate two algorithms, and the algorithms are switched on time for processing. Another possibility is that two cards are interconnected and used to accelerate a single algorithm. A third possibility is to exploit the coarse grain parallelism inherent in a certain algorithm, and use more than one card to perform the same function, producing data results in parallel. The final architecture will be established as the ALOA system is developed.

The hardware design process utilized by the ASDP group uses a hardware description language, VHDL, to express the design in a form closer to the programming language used in the microprocessor based host environment. The ASDP group has also developed an incremental design, integration and test procedure that facilitates the application development process. The development of adaptive computing applications is still in art form. Several research groups are investigating and developing automated processes and tools for the development of applications. The process today is a combination of software and hardware engineering, and thus requires a team of engineers that comprises a wide set of skills.

Conclusion

We presented the plan for the construction of a prototype named Adaptive Level One Accelerator (ALOA). The system intends to process and produce the MODIS level one products at near real time data rates. It is expected that by utilizing a new data processing technology, adaptive computing, the ALOA will improve the cost performance ratio by two orders of magnitude. The prototype also serve as a proof of concept for the utilization of adaptive computing to compose a system that has the performance of hardware oriented architectures and the flexibility and maintanability of software systems.

Finally, the successful accomplishment of a real-time direct readout system for high data rate instruments at low cost will enable a new era of remote sensing products. The ALOA prototype serves as the path finder for this new era.

References

[1]Nasa Goddard Space Flight Center; Direct Access System User's Guide for the EOS-AM Spacecraft; ICD-107; IS20008696; November 20, 1998.

[2]Santa Barbara Research Center; MODIS Command, Telemetry, Science and Engineering Description; Doc. Number 151840; May 14, 1997.

[3]MODIS Science Data Support Team; MODIS Level 1A Earth Location: Algorithm Theoretical Basis Document – Version 3.0; August 26, 1997

[4]MODIS Characterization Support Team; Draft of the MODIS Level 1B Algorithm Theoretical Basis Document – Version 2.0; February 13, 1997

[5]Guenther B., Godden G. D., Xiong X., Knight E. J., Qiu S., Montgomery H., Hopkins M. M., Khayat M. G., Hao Z.; "Prelaunch Algorithm and Data Format for the Level 1 Calibration Products for the EOS-AM1 Moderate Resolution Imaging Spectroradiometer (MODIS); IEEE Transactions on Geoscience and Remote Sensing, vol. 36, No. 4, July 1998

[6]Figueiredo M., Gloster C.; "Implementation of a Probabilistic Neural Network for Multi-Spectral Image Classification on a FPGA Custom Computing Machine", 5th Brazilian Symposium on Neural Networks – Belo Horizonte, Brazil, December 9-11, 1998

[7]Graessle T., Figueiredo M.; "Application of Adaptive Computing in Satellite Telemetry Processing"; 34th International Telemetering Conference - San Diego, California, October 26-29, 1998

[8]Figueiredo M., Graessle T., Winiecki K., Greenspan, D.; "MODIS Reflective Calibration Acceleration – ASDP Fall 1997 Report"; http://fpga.gsfc.nasa.gov/asdp/fall97rep/fall97rep.htm